



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 087 576 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
28.03.2001 Bulletin 2001/13

(51) Int Cl. 7: H04L 12/56

(21) Application number: 00650118.3

(22) Date of filing: 28.08.2000

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SEDesignated Extension States:
AL LT LV MK RO SI

(30) Priority: 31.08.1999 US 387036

(71) Applicant: Nortel Networks Limited
Montreal, Quebec H2Y 3Y4 (CA)(72) Inventor: Hsu, Ivy P.
Danville, California 94506 (US)(74) Representative: Coyle, Philip Aidan et al
F. R. KELLY & CO.
27 Clyde Road
Ballsbridge
Dublin 4 (IE)

(54) Constraint-based route selection using biased cost

(57) The present invention is a method and apparatus for selecting a route for a flow from a plurality of network paths connecting a source to a destination. The method comprises: (a) determining cumulative costs for a plurality of candidate paths from the network paths us-

ing a cost bias which is dynamically calculated based on at least one of a flow attribute and a path attribute; and (b) selecting an optimal path having a minimum of the cumulative costs. The optimal path corresponds to the selected route.

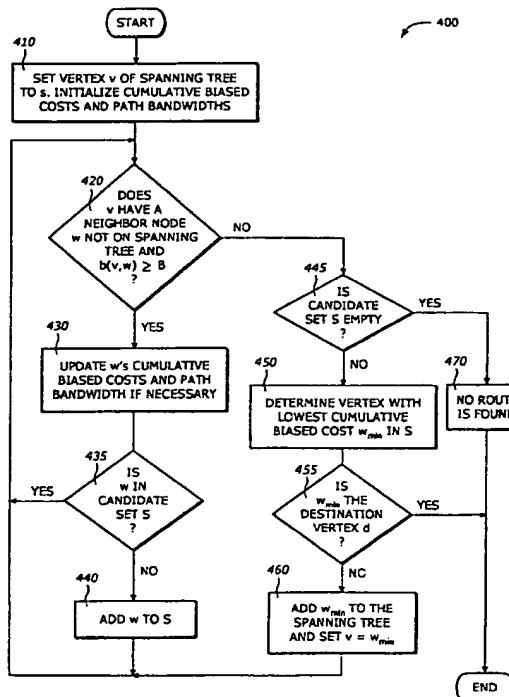


FIG. 4

Description	BACKGROUND
1. Field of the Invention	The cumulative cost for a candidate node in a candidate path includes a current cumulative cost and a biased static cost. The biased static cost includes a static cost of the link biased by a bias value. The bias value is a function of the flow priority, the bandwidth demand, the link bandwidth and a maximum available link bandwidth. The bandwidth requirement is to route selection with bandwidth width requirement is to modify the well-known Dijkstra technique so that only links that meet the bandwidth requirement are considered in the Dijkstra procedure. The result is a minimum-cost route over links that meet the bandwidth requirement.
2. Description of Related Art	[0002] One common approach in route selection with bandwidth width requirement is to route selection with bandwidth width requirement is to modify the well-known Dijkstra technique so that only links that meet the bandwidth requirement are considered in the Dijkstra procedure. The result is a minimum-cost route over links that meet the bandwidth requirement.
3. Brief Description of the Drawings	[0003] For example, in a packet forwarding known as Multi-Protocol Label Switching (MPLS), to accommodate different priorities, the label switched paths (LSPs) of different priorities, the label selector may first apply the modified Dijkstra routing selection based on a given LSP's bandwidth demand. If a feasible path is not found, the route selector may preempt an existing LSP of lower priority from a connection to a new LSP. A pre-empted LSP will have to be re-established on an alternative route.
4. Detailed Description of the Invention	[0004] This approach has a number of shortcomings. First, it is difficult to make an intelligent decision of which LSP to preempt. The traditional technique is to consider all candidate LSPs on all possible links, evaluate the minimum cost impact, and select the LSP that has the minimum cost of preempting, and select the LSP that has the minimum cost impact. This is computationally expensive and may fail to be preempted, resulting in loss of connection. Third, the approach disrupts the LSP's traffic and decreases service quality due to effort in pre-empting and re-routing. Last, the approach may lead to traffic contention a network because some high priority LSPs may be placed on very long routes.
5. Summary	[0005] Therefore there is a need in the technology to provide a simple and efficient method to select routes in a system of networks.
6. Brief Description of the Drawings	[0006] The present invention is a method and apparatus for selecting a route for a flow from a plurality of routes for selecting paths. Cumulative costs for a plurality of candidate paths from the network paths are determined using network paths. Cumulative costs for a flow from a plurality of routes for selecting a route for a flow according to one embodiment of the invention.
7. Detailed Description of the Invention	[0007] According to one embodiment of the present invention, the flow attribute includes a flow priority and a selected route. The optimal path corresponds to the cumulative costs. The selected path has a minimum of the cumulative costs which is dynamically calculated based on at least one of a flow attribute and a path attribute. An optimal path is then selected which has a minimum of the cumulative costs. The optimal path corresponds to the selected route.
8. Brief Description of the Drawings	[0008] The present invention relates to computer networks. In particular, the invention relates to route selection. In particular, the invention relates to route selection.
9. Detailed Description of the Invention	[0009] Other aspects and features of the present invention will become apparent from the following detailed description of the present invention in which:
10. Brief Description of the Drawings	[0010] The features and advantages of the present invention will become apparent from the following detailed description of the present invention in which:
11. Detailed Description of the Invention	[0011] Figure 1A is a diagram illustrating a system having a central server with a biased cost route selector in which local biased cost route selection can be practiced. In which one embodiment of the invention can be practiced.
12. Brief Description of the Drawings	[0012] Figure 1B is a diagram illustrating a system with local biased cost route selection in which one embodiment of the invention can be practiced.
13. Detailed Description of the Invention	[0013] Figure 1C is a diagram illustrating a system with inter-area routing in which one embodiment of the invention can be practiced.
14. Brief Description of the Drawings	[0014] Figure 1D is a diagram illustrating a computer system in which one embodiment of the invention can be practiced.
15. Detailed Description of the Invention	[0015] Figure 2 is a diagram illustrating a typical router according to one embodiment of the invention.
16. Brief Description of the Drawings	[0016] Figure 3 is a diagram illustrating a graph representing a network topology according to one embodiment of the invention.
17. Detailed Description of the Invention	[0017] Figure 4 is a flowchart illustrating a process to select route using biased cost according to one embodiment of the invention.
18. Brief Description of the Drawings	[0018] Figure 5 is a flowchart illustrating a process to initialize the graph parameters according to one embodiment of the invention.
19. Detailed Description of the Invention	[0019] Figure 6 is a flowchart illustrating a process to add candidate vertices according to one embodiment of the invention.
20. Brief Description of the Drawings	[0020] Figure 7 is a flowchart illustrating a process to provide admission control according to one embodiment of the invention.
21. Detailed Description of the Invention	[0021] Figure 8 is a flowchart illustrating a process to for optimistic inter-area routing according to one embodiment of the invention.
22. Brief Description of the Drawings	[0022] Figure 8B is a flowchart illustrating a process for conservative inter-area routing according to one embodiment of the invention.

bodiment of the invention.

[0023] Figure 9A is a diagram illustrating an exponential cost bias function according to one embodiment of the invention.

[0024] Figure 9B is a diagram illustrating a linear combination cost bias function according to one embodiment of the invention.

[0025] Figure 9C is a diagram illustrating a logarithmic cost bias function according to one embodiment of the invention.

DESCRIPTION

[0026] The present invention relates to a method and apparatus for selecting routes in computer communication networks. The technique determines cumulative costs of candidate paths using a biased cost function and selects the optimal path based on the minimum cumulative costs. The biased cost function allows the selection be made intelligently and dynamically to accommodate the traffic requirements and bandwidth availability.

[0027] In the following description, for purposes of explanation, numerous details are set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these specific details are not required in order to practice the present invention. In other instances, well known electrical structures and circuits are shown in block diagram form in order not to obscure the present invention.

[0028] The present invention is a constraint-based route selection technique that supports establishing Multi-protocol Label Switching (MPLS) label switched paths through explicit routing. The invention relates to applications where the constraints associated with a flow include a bandwidth requirement and a priority. The bandwidth requirement provides that a flow should be routed along a path with sufficient available bandwidth on each link. The priority attribute differentiates flows by the relative likelihood of being blocked by the network due to insufficient resources, so that flows of the higher priority class should expect higher service availability.

[0029] The present invention introduces a number of modifications and improvements to the standard Dijkstra calculation to accommodate these two constraints. The bandwidth requirement is met by considering only links with sufficient bandwidths in the iterative procedure within Dijkstra. In determining the cumulative cost of each potential route, the technique calculates a cost bias factor for each link, which is a function of the link's bandwidth availability and the priority and bandwidth of the given flow. Thus, the cost of a link is the product of its static cost from link state advertisements and this bias factor. The technique selects a route with sufficient bandwidth that minimizes the cumulative biased cost.

[0030] The objective of the bias factor is to make highly congested links less desirable for new flows. Further-

more, this sensitivity to link congestion is greater for lower priority classes, so that low priority flows with less costly alternate paths will avoid these congested links, leaving the remaining bandwidths for high priority flows.

5 The intent is to keep high priority flows on their direct paths and place low priority flows on longer alternate paths if necessary. Thus, this technique is referred to as a multi-class technique using cost bias.

[0031] The route selection process can be implemented either centrally in a network server or distributedly on each label edge router (LER). The technique is designed to be simple enough to be implementable on routers and retains the same order of complexity as the standard Dijkstra technique.

10 [0032] Figure 1A is a diagram illustrating a system 100A having a central server with a biased cost route selector in which one embodiment of the invention can be practiced. The system 100A includes a central server 110, routers 120A1 to 120AK, and networks 122A1 to 122AK. The central server 110 includes a biased cost route selector 115A.

15 [0033] The central server 110 acts as a centralized server for the entire system of networks. The central server 110 forwards the route decision to each of the routers 120A1 to 120AK. The biased cost route selector 115A provides the central server 110 the route selection decisions.

20 [0034] The routers 120A1 to 120AK provides connectivity among the networks 122A1 to 122AK. The networks 122A1 to 122AK are any computer networks that provide network communications.

25 [0035] Figure 1B is a diagram illustrating a system 100B with local biased cost route selectors in which one embodiment of the invention can be practiced. The system 100B includes routers 120Bi and networks 122Bi, where i = 1, ..., N.

30 [0036] The routers 120Bi route the traffic flows in the system to the associated networks 122Bi. Each of the routers 120Bi has a biased cost route selector (BCRS) 115Bi. The BCRS 115Bi selects the routes based on an optimal biased cost metric as will be explained later. In this system, the BCRS 115Bi performs route selection in a distributed manner.

35 [0037] Figure 1C is a diagram illustrating a system 100C with inter-area routing in which one embodiment of the invention can be practiced. The system 100C includes an area 102 and a backbone 103.

40 [0038] The area 102 includes routers 120C1, 120C2, 120C3, 120C4, and 120C5. The routers 120C1, 120C2, 120C3, 120C4, and 120C5 are coupled to the networks 122C1, 122C2, 122C3, 122C4, and 122C5, respectively. The backbone 103 includes routers 120C6 and 120C7 coupled to networks 122C6 and 122C7, respectively. Each of the routers has a biased cost route selector. The router 120C1 is an area router as well as a backbone router. It advertises the summary information about the networks in the area 102 into the backbone 103 and vice versa. Suppose the summary LSAs are

cost to that destination. Therefore, if a network has all router links of cost 1, the cost metric becomes equivalent to hop count and the least-cost path is simply the shortest-hop path.

[0050] Each router in the network maintains an image of the topology (V, E) and $c(v, w)$ for all $(v, w) \in E$ through standard OSPF link state advertisements. QoS extensions provide advertisements of $b(v, w)$ for all $(v, w) \in E$. Note that in general, dynamic events such as link additions, outages, and changes in available bandwidths propagate throughout the network after some finite delay, so that at any time the topology information may be out of date.

[0051] In the exemplary graph of Figure 3:

$$V = \{v_1, v_2, v_3, v_4, v_5, v_6\}$$

$$E = \{(v_1, v_3), (v_3, v_1), (v_1, v_4), (v_2, v_4), (v_4, v_2), (v_3, v_5), (v_2, v_5), (v_5, v_2), (v_2, v_6), (v_6, v_2), (v_6, v_1)\}$$

[0052] The links (v_1, v_3) , (v_3, v_1) , (v_1, v_4) , (v_2, v_4) , (v_4, v_2) , (v_3, v_5) , (v_2, v_5) , (v_5, v_2) , (v_2, v_6) , (v_6, v_2) , and (v_6, v_1) have the current available bandwidths $b(v_1, v_3)$, $b(v_3, v_1)$, $b(v_1, v_4)$, $b(v_2, v_4)$, $b(v_4, v_2)$, $b(v_3, v_5)$, $b(v_2, v_5)$, $b(v_5, v_2)$, $b(v_2, v_6)$, $b(v_6, v_2)$, and $b(v_6, v_1)$, and cost metrics $c(v_1, v_3)$, $c(v_3, v_1)$, $c(v_1, v_4)$, $c(v_2, v_4)$, $c(v_4, v_2)$, $c(v_3, v_5)$, $c(v_2, v_5)$, $c(v_5, v_2)$, $c(v_2, v_6)$, $c(v_6, v_2)$, and $c(v_6, v_1)$, respectively.

BIASED COST ROUTE SELECTION:

[0053] The objective of the route selection is to find a feasible explicit path for a CR-LSP from source s to destination d with a bandwidth requirement B . The CR-LSP has an associated priority $m \in M$. A numerically smaller m represents a higher priority.

[0054] m corresponds to the Setup Priority which represents the LSP's availability attribute. That is, the likelihood of the LSP being rejected due to lack of resources should be lower for lower m . In addition, the Holding Priority in MPLS represents an LSP's resilience attribute in being preempted. It is defined such that LSPs with high availability always have high resilience, while LSPs with low availability can also be promoted to have higher resilience. The route placement takes the Setup Priority m into account. It does not explicitly use the Holding Priority. Furthermore, the routing priority may or may not be associated with any service scheduling priority (as in Differentiated Services).

[0055] The biased cost route selector introduces a cost bias to the static OSPF cost metric $c(v, w)$ associated with link (v, w) . This cost bias is calculated dynamically by the route selection process handling the given flow (at either the ingress LER or a centralized server). Again, note that this is different from having each router calculate dynamic cost metrics for its outgoing links and advertising them in the OSPF link state advertisements.

[0056] There are a number of advantages to this ap-

proach: (1) since hop-by-hop routing is still based on the advertised cost metrics, the biased cost route selection technique leaves the hop-by-hop routing system intact; and (2) the cost bias in this case can depend on not only the current state of the link resource, $b(v, w)$, but also the characteristics of the flow to be placed (i.e., its bandwidth demand B and priority m).

[0057] Furthermore, define $b_{max} = \max b(c, w)$ for all $(v, w) \in E$, b_{max} is the maximum available link bandwidth in the network at the current state. It is used to obtain a normalized link availability $(- b(v, w)/b_{max})$. This normalized link availability is used instead of each link's own relative availability ($= b(v, w)/\text{physical capacity}$) for a number of reasons: (1) It better reflects availability in network with different physical link speeds, e.g., a new OC-3 has a greater availability than a new DS-3. (2) A link's physical capacity may not be advertised in the QoS LSAs. (3) The ratio tracks the relative utilization of links as network loading evolves.

[0058] The cost bias for link (v, w) is therefore denoted as $p(m, B, b(v, w), b_{max})$ to indicate its dependence on these four parameters. The biased cost route selector modifies the shortest-widest feasible path algorithm by replacing the static cost $c(v, w)$ used in the Dijkstra process by a biased cost $c(v, w) * p(m, B, b(v, w), b_{max})$. In other words, it selects the path with the minimum cumulative biased cost $c(v, w) * p(m, B, b(v, w), b_{max})$ for all (v, w) belonging to r , subject to the feasibility constraint $b(v, w) B$, for all (v, w) belonging to r . The path bandwidth serves as a tie-breaker when there are more than one path with the same cumulative biased cost.

[0059] Denote

v = vertex under consideration;

S = the set of candidate vertices to be added to the spanning tree;

$L(x)$ = cumulative biased cost from source s to vertex x ; and

$B(x)$ = path bandwidth from source s to vertex x .

[0060] The biased cost route selection can be performed in the following steps.

Step 1: Initialization

[0061] Set $v = s$. s is the root of the spanning tree.

[0062] Let $L(v) = 0$ and $B(v) = \infty$. $L(x) = \infty$ and $B(x) = 0$ for all other vertices.

[0063] Step 2: Update metrics and add candidates to the set S

[0064] Perform the following updates for each vertex w that (1) is not already on the spanning tree, (2) has an edge $(v, w) \in E$, and (3) has sufficient bandwidth $b(v, w) \geq B$:

if $L(w) < L(v) + c(v, w) * p(m, B, b(v, w), b_{max})$, a lower cost feasible path from s to w is found through v .

420.

[0076] Once all neighbors of v have been examined, the process 400 determines if the candidate set S is empty (Block 445). If the candidate set S is empty, the process 400 determines that no route is found and inform the appropriate control (Block 470) and is then terminated. If the candidate set S is not empty, the process 400 determines the vertex in the candidate S with the lowest cumulative biased cost w_{min} (Block 450). Then the process 400 determines if w_{min} is the destination vertex d . If it is, the best route has been found and the process 400 is terminated. Otherwise, the process 400 adds w_{min} to the spanning tree and set $v = w_{min}$ (Block 460). Next, the process 400 returns to Block 420.

[0077] Figure 5 is a flowchart illustrating a process 410 to initialize the graph parameters according to one embodiment of the invention.

[0078] Upon START, the process 410 initializes the set S to an empty set (Block 510). Then the process 410 sets all the cumulative biased costs $L(x)$ to infinity and sets all the path bandwidths $B(x)$ to zero (Block 520).

[0079] Then the process 410 selects the source vertex s as the first candidate vertex v (Block 530). Next, the process 410 sets the cumulative biased cost $L(v)$ to zero and the path bandwidth $B(v)$ to infinity (Block 540). Then the process 410 is terminated.

[0080] Figure 6 is a flowchart illustrating a process 430 to update the cumulative biased cost according to one embodiment of the invention.

[0081] Upon START, the process 430 computes the temporary value $D(w) = L(v) + c(v,w)*p \{m, B, b(v,w), b_{max}\}$ (Block 610) where $c(v,w)$ is the static cost of link (v,w) and $p(m, B, b(v,w), b_{max})$ is the biased cost function.

[0082] Then the process 430 determines if the cumulative cost $D(w)$ is less than $L(w)$ (Block 620). If not, the process 430 determines if $D(w)$ is equal to $L(w)$ and $B(w)$ is less than $\min(B(v), b(v,w))$ (Block 640). If no, no updating is necessary and the process 430 is terminated. If $D(w)$ is equal to $L(w)$ and $B(w)$ is less than $\min(B(w), b(v,w))$, the process 430 updates $B(w)$ by setting $B(w)$ to $\min \{B(v), b(v,w)\}$ (Block 650) and is then terminated.

[0083] If the cumulative cost $D(w)$ is less than $L(w)$, the process 430 updates $L(w)$ and $B(w)$ by setting $L(w)$ to $D(w)$ and $B(w)$ to $\min \{B(v), b(v,w)\}$ (Block 630). Then the process 430 is terminated.

Incorporating Admission Control with Routing:

[0084] In either the multi-class biased-cost route selection or the shortest-widest feasible path selection, a flow is routed so long as a feasible route exists. This may not always be desirable, as the cost in resource consumed (and hence future blocking) may exceed the gain in placing the flow. This phenomenon has been observed in circuit-switching networks in the past. In circuit-switching networks, trunk reservation provides an

effective protection against the undesirable condition in which most circuits are routed along indirect paths. Trunk reservation can be viewed as a form of utilization-sensitive admission control: when the amount of bandwidth used by indirect circuits on a given link exceeds a threshold, no additional indirect circuits are allowed on the link.

[0085] When a route is selected by the shortest-widest feasible path algorithm, the cumulative path cost reflects the length of the route (and administrative weight of the links if cost $\neq 1$), but not the utilization. With the biased cost algorithm, a high cumulative cost can reflect high utilization, long route, or a combination of the two. The network can therefore elect to impose an admission control rule on each flow priority as follows:

[0086] Let r^* be the least biased cost route found for a given flow (s, d, B, m) . Reject the flow if

$$20 \quad c(v,w)*p(m, B, b(v,w), b_{max}) > T(m),$$

where $T(m)$ is the admission threshold for flows of priority m .

[0087] In other words, $T(m)$ reflects the benefit obtained by the network for routing a flow of priority m . The route selection algorithm should identify the route that maximizes the net benefit ($T(m) - \text{cost}$) if one exists, and reject the flow if the routing cost exceeds the benefit.

[0088] For example, by setting $T(0) = \infty$ and $T(m) < \infty$ for other $m \in M$, the network rejects lower-priority flows from routes with high-utilization links and/or long hops, which helps preserving the retraining capacity for use by highest-priority flows and reducing their rejection probability. We note that this is achieved at the expense of increased rejection probability for lower priorities. This is analogous to the effect of discarding schemes such as ATM Cell Loss Priority (CLP) threshold and Random Early Drop with In/Out bit (RIO) which preserve buffer resources for higher priority traffic at the expense of discarding lower priority traffic early.

[0089] Figure 7 is a flowchart illustrating a process 700 to provide admission control according to one embodiment of the invention.

[0090] Upon START, the process 700 determines the least biased cost route r^* for a given flow (s, d, B, m) (Block 710). This is carried out as illustrated in Figure 4. Then the process 700 determines if the sum of the biased cost for the priority m is greater than the admission threshold (Block 720). If no, the process 700 accept the route r^* (Block 740) and is then terminated. If yes, the process 700 rejects the route r^* (Block 730) and perhaps attempts to find another route. Then the process 700 is terminated.

55 Inter-Area Routing:

[0091] OSPF introduces the concept of area to improve Autonomous System (AS) scalability through to-

[0098]	In the second, or conservative, approach, the route selection determines a cost bias for each ABR based on its available path bandwidth to the destination, using the same formula as link cost bias. This bias is then multiplied with the path cost from the ABR to that destination, and the route based cost from the ABR to the destination, which is why the same bias value can be applied to the entire inter-area segment. Then the process 800 is terminated.	1. Consider all ABRs with sufficient bandwidth to the destination.	2. Find a shortest-widest feasible path within the area destination.	3. If there are multiple eligible ABRs, select the one source to destination as a tie-breaker.	40 [0094] Cost bias can be incorporated in this frame work in one of two ways: an optimistic method and a conservative method.
[0099]	Figure 8B is a flowchart illustrating a process 800 for conserving inter-area routing according to one area segment.	1. Consider all ABRs with sufficient bandwidth to the destination.	2. Find a shortest-widest feasible path within the area destination.	3. If there are multiple eligible ABRs, select the one source to destination as a tie-breaker.	30 [0095] In the first, or optimistic, approach, the route based cost to destination is made based on the sum of the path cost from the ABR and the cumulative cost from the ABR source to the destination. In Step 2 above, using minimum bias based cost to each ABR in Step 3, the static cost for the intra-area segment is made based on the sum of the biased cost for the intra-area segment and the static cost for the inter-area segment. In other words, no bias is applied to the area segment.
[0100]	Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	40 [0095] In the first, or optimistic, approach, the route based cost to destination is made based on the sum of the path cost from the ABR and the cumulative cost from the ABR source to the destination. In Step 2 above, using minimum bias based cost to each ABR in Step 3, the static cost for the intra-area segment is made based on the sum of the biased cost for the intra-area segment and the static cost for the inter-area segment. In other words, no bias is applied to the area segment.	45 [0096] Figure 8A is a flowchart illustrating a process 800 for conserving inter-area routing according to one area segment.	45 [0096] Figure 8A is a flowchart illustrating a process 800 for conserving inter-area routing according to one area segment.	45 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.
[0101]	In either solution, some optimization potential is lost due to the detail-hiding property of topology hierarchy. These are the necessary tradeoffs to achieve scalability. Furthermore, the partition into intra- and inter-area segments means that not all potential source destinations find the feasible intra-area path width.	50 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	50 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	50 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	50 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.
[0102]	This provides one motivation for a centralized route selection server that monitors and maintains global link state information for the entire ABRs and all areas. This device may contain a library of route descriptions (AS), and generates explicit routes for individual item (AS), and generates explicit routes for individual item (AS).	55 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	55 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	55 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	55 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.
[0103]	A centralized route selection server is a device that maintains a topology database of the entire ABRs and all areas. This device may contain a library of route descriptions (AS), and generates explicit routes for individual item (AS), and generates explicit routes for individual item (AS).	60 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	60 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	60 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.	60 [0097] Upon START, the process 800 considers all area border routers (ABRs) with sufficient bandwidth to the destination.

bandwidth/priority/cost, delay/cost, diversity, load balancing). For each given CR-LSP, the route selection server may consult the network policy repertoire for the suitable objective and select the algorithm to run. This route selection server may work as an advisory tool for network operators in traffic placement, or, in an automated environment, generate contents of CR-LSP TLVs and forward them to the source LERs.

[0104] A centralized server can help select inter-area paths that are optimized for an AS. We should also note that since the server has the complete topological information, the explicit route can be requested with all intermediate nodes specified, instead of loosely routed for the segment external to the area of the source LER. Furthermore, since the centralized server has the complete bandwidth availability information within an AS, for the multi-class technique it can apply the correct cost bias for each link between the source and the destination. This avoids the use of approximations for links outside of the source LER's.

[0105] In addition, certain optimizations to the multi-class technique may not be viable in a distributed implementation and only suitable in a centralized server paradigm. One such case is parameter optimization. The proposed cost bias functions contain one or more parameters associated with the priority classes.

[0106] Parameter optimization is useful when the network topology and offered traffic have stationary or slow-varying characteristics, such that there is a set of parameters optimizing the routing policy's reward. If the network and offered load are time-varying, as may be expected in practical networks, the benefit of training and tuning the parameters may not warrant the computational overhead.

[0107] A centralized server paradigm is also more suitable to support multiple route selection algorithms that are policy-driven. Unlike a distributed implementation, introducing a new algorithm or a modification to reflect a policy change does not require software upgrade on all LERs.

COST BIAS FUNCTION:

Desirable Properties of the Cost Bias Function:

[0108] The cost bias function $p(m, B, b(v,w), bmax)$ biases the static cost of link (v, w) to reflect its bandwidth availability, such that the selected path has the minimum cumulative biased cost from s to d . In the event of multi-paths with equal biased cost, the path bandwidth serves as a tie-breaker as before. Furthermore, flows with different priorities are subject to different degrees of bias to reflect their sensitivities toward congested links. The cost bias function $p(m, B, b(v,w), bmax)$ should have the following properties:

1. $p(m, B, b(v,w), bmax) < p(n, B, b(v,w), bmax)$ for $m < n$ where m and $n \in M$. The cost to place a given

bandwidth B on a given link with $b(v, w)$ should be higher for lower priority (higher numerical value) flows.

5 2. $dp(m, B, b(v,w), bmax)/ db(v,w) < 0$. That is, the larger is the available capacity, the less the cost bias.

10 3. $dp(m, B, b(v,w), bmax)/ dB > 0$. That is, the larger is the bandwidth demand, the higher the cost bias.

15 4. $dp(m, B, b(v,w), bmax)/ db(v,w) > dp(n, B, b(v,w), bmax)/ db(v,w)$ and $dp(m, B, b(v,w), bmax)/ dB < dp(n, B, b(v,w), bmax)/ dB$ for $m < n$ where m and $n \in M$. Since n is a lower priority, it should have a steeper cost bias curve. In other words, lower priority flows should be more sensitive toward a congested link.

20 5. $1 p(m, B, b(v,w), bmax) < 0$ for $m \in M$, $0 B bmax$, and $0 b(v,w) bmax$. The lower bound is not necessary, but is desirable so that the biased cost is always greater than or equal to the static cost.

25 Special Cases:

[0109] Furthermore, based on these properties of the cost bias function, when the parameters are reduced to a number of special cases, the biased cost route selector should have the following characteristics:

[0110] Single priority class: The biased cost route selector is also applicable in classless routing. In this case, $p(m, B, b(v,w), bmax)$ is reduced to $p(B, b(v,w), bmax)$. The resulting route selection algorithm is one that biases against choosing highly congested links and loads traffic among multiple paths between a source-destination pair by trading off hop count with link congestion.

[0111] Equal-loading links: When $b(v, w) = bmax$ for all $(v, w) \in E$, the technique should select the same path as the shortest-widest feasible path technique. This is because with any given pair of m and B , all links are biased equally by a constant $p(m, B)$, i.e., $p(m, B, b(v, w), bmax) = p(m, B)$. Thus

$$45 c(v,w) * p(m, B, b(v,w), bmax) = p(m, B) * c(v,w)$$

for all (v,w) belonging to r and r belonging to R .

[0112] Only the cumulative static costs affect the route selection and the multi-class algorithm behaves the same as the shortest-widest feasible path.

[0113] Extending this further, if all links have the same physical bandwidth and the link utilization within the network remains uniform, the routing behaviors of the two methods should remain close. However, when the traffic loading to the network is biased toward certain sets of source-destination pairs, which arises in most practical scenarios, the multi-class technique attempts to offer a

Claims	Cost Bias Functions:	greater diversification on the multi-paths.
1. A method to select a route for a flow from a plurality of network paths, comprising:	[0144] A normalized link availability (v, w) is used in the cost bias functions $p(m, (v, w)) = \exp \{am \cdot (1 - (v, w))\}$, where $1 \leq am < an$ for $m < n$, m and $n \in M$.	[0115] Three families of cost bias functions that satisfy the desirable properties above are: - Exponential: $p(m, (v, w)) = \exp \{am \cdot (1 - (v, w))\}$, where $1 \leq am < an$ for $m < n$, m and $n \in M$. - Linear combination: $p(m, (v, w)) = 1 \cdot (v, w) - 2am + 1 + 2 \cdot (v, w) - am^2$, where $0 \leq am < an$ for $m < n$, m and $n \in M$. - Logarithmic: $p(m, (v, w)) = (\log(v, w)) + 1$ - arm, where $0 < am < an$ for $m < n$, m and $n \in M$.
2. The method of claim 1 wherein the flow attribute includes a flow priority and a bandwidth demand.	[0116] Figure 9A is a diagram illustrating an exponential cost bias function according to one embodiment of the invention.	[0117] There are four curves 910, 912, 914, and 916 corresponding to four classes of flows. These four curves all have $1 = 2 = 0.5$. The parameters for the am, curves 920, 922, 924, and 926 are $1 = 0.1$, $1, 1 = 0.2 = 0.1$, $2, 1 = 0.2 = 0.1$, and $3, 1 = 0.3 = 0.2$. The minimum of the optimal path having cumulative costs is lower than one minimum of the cumulative costs, selecting the optimal path having the minimum path bandwidth.
3. The method of claim 2 wherein the path attribute includes a link bandwidth and a maximum available link bandwidth.	[0118] Figure 9B is a diagram illustrating a linear combination cost bias function according to one embodiment of the invention.	[0119] There are four curves 920, 922, 924, and 926 corresponding to four classes of flows. These four curves all have $1 = 2 = 0.5$. The parameters for the am, curves 930, 932, 934, and 936 are $1 = 0.1$, $1, 1 = 0.2 = 0.1$, $2, 1 = 0.2 = 0.1$, and $3, 1 = 0.3 = 0.4$, respectively.
4. The method of claim 3 further comprising:	[0120] Figure 9C is a diagram illustrating a logarithmic cost bias function according to one embodiment of the invention.	[0121] There are four curves 930, 932, 934, and 936 corresponding to four classes of flows. The curves 930, 932, 934, and 936 have the parameter values $am = 0.5$, $1 = 1, 1 = 0.2 = 0.1$, $2, 1 = 0.2 = 0.1$, and $3, 1 = 0.3 = 0.2$. The minimum of the optimal path having a cumulative cost includes a static cost of the link biased by a bias value.
5. The method of claim 4 wherein determining the cumulative costs comprises:	[0122] While this invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments, will be apparent to persons skilled in the art to which the invention pertains are deemed to lie within the spirit and scope of the invention.	[0122] The method of claim 7 wherein an increase in link bandwidth, the link bandwidth, and the maximum available link function of the flow priority, the bandwidth demand, the link bandwidth, and the bias value is a priority increases the bias value.
6. The method of claim 4 wherein selecting the optimal path comprises:	[0123] Figure 9D is a diagram illustrating a logarithmic cost bias function according to one embodiment of the invention.	[0123] The method of claim 7 wherein a decrease in link bandwidth, the link bandwidth, and the bias value is a priority decreases the bias value.
7. The method of claim 6 wherein the candidate bidased cumulative cost includes a current cumulative cost and a biased static cost of the link biased by a bias value.	[0124] Figure 9E is a diagram illustrating a logarithmic cost bias function according to one embodiment of the invention.	[0124] The method of claim 7 wherein a decrease in link bandwidth, the link bandwidth, and the bias value is a priority decreases the bias value.
8. The method of claim 7 wherein the bias value is a function of the flow priority, the bandwidth demand, the link bandwidth, and the bias value.	[0125] Figure 9F is a diagram illustrating a logarithmic cost bias function according to one embodiment of the invention.	[0125] The method of claim 7 wherein a decrease in link bandwidth, the link bandwidth, and the bias value is a priority decreases the bias value.
9. The method of claim 7 wherein a decrease in flow priority increases the bias value.	[0126] Figure 9G is a diagram illustrating a logarithmic cost bias function according to one embodiment of the invention.	[0126] The method of claim 7 wherein a decrease in link bandwidth, the link bandwidth, and the bias value is a priority decreases the bias value.

bandwidth decreases the bias value.

11. The method of claim 7 wherein an increase in bandwidth demand increases the bias value.

5

12. The method of claim 7 wherein the bias value is greater than or equal to 1.

13. The method of claim 6 further comprising:

rejecting the selected optimal path if the minimum of the cumulative costs exceeds an admission threshold.

10

14. The method of claim 6 wherein the plurality of network paths include an intra-area segment and an inter-area segment.

15

15. The method of claim 14 wherein the cumulative costs in the inter-area segment exclude the cost bias.

20

16. The method of claim 14 wherein the inter-area segment has the same cost bias for the candidate paths.

25

17. A computer program product comprising:

a computer usable medium having computer program code embodied therein to select a route for a flow from a plurality of network paths, the computer program product having:
30 computer readable program code for determining cumulative costs for a plurality of candidate paths from the network paths using a cost bias, the cost bias being dynamically calculated based on at least one of a flow attribute and a path attribute; and
35 computer readable program code for selecting an optimal path having a minimum of the cumulative costs, the optimal path corresponding to the selected route.

40

18. A system comprising:

a processor; and
45 a memory coupled to the processor, the memory containing program code to select a route for a flow from a plurality of network paths, the program code when executed causing the processor to determine cumulative costs for the network paths using a cost bias, the cost bias being dynamically calculated based on at least one of a flow attribute and a path attribute, and select an optimal path having a minimum of the cumulative costs, the optimal path correspond-
50 ing to the route.

55

100A

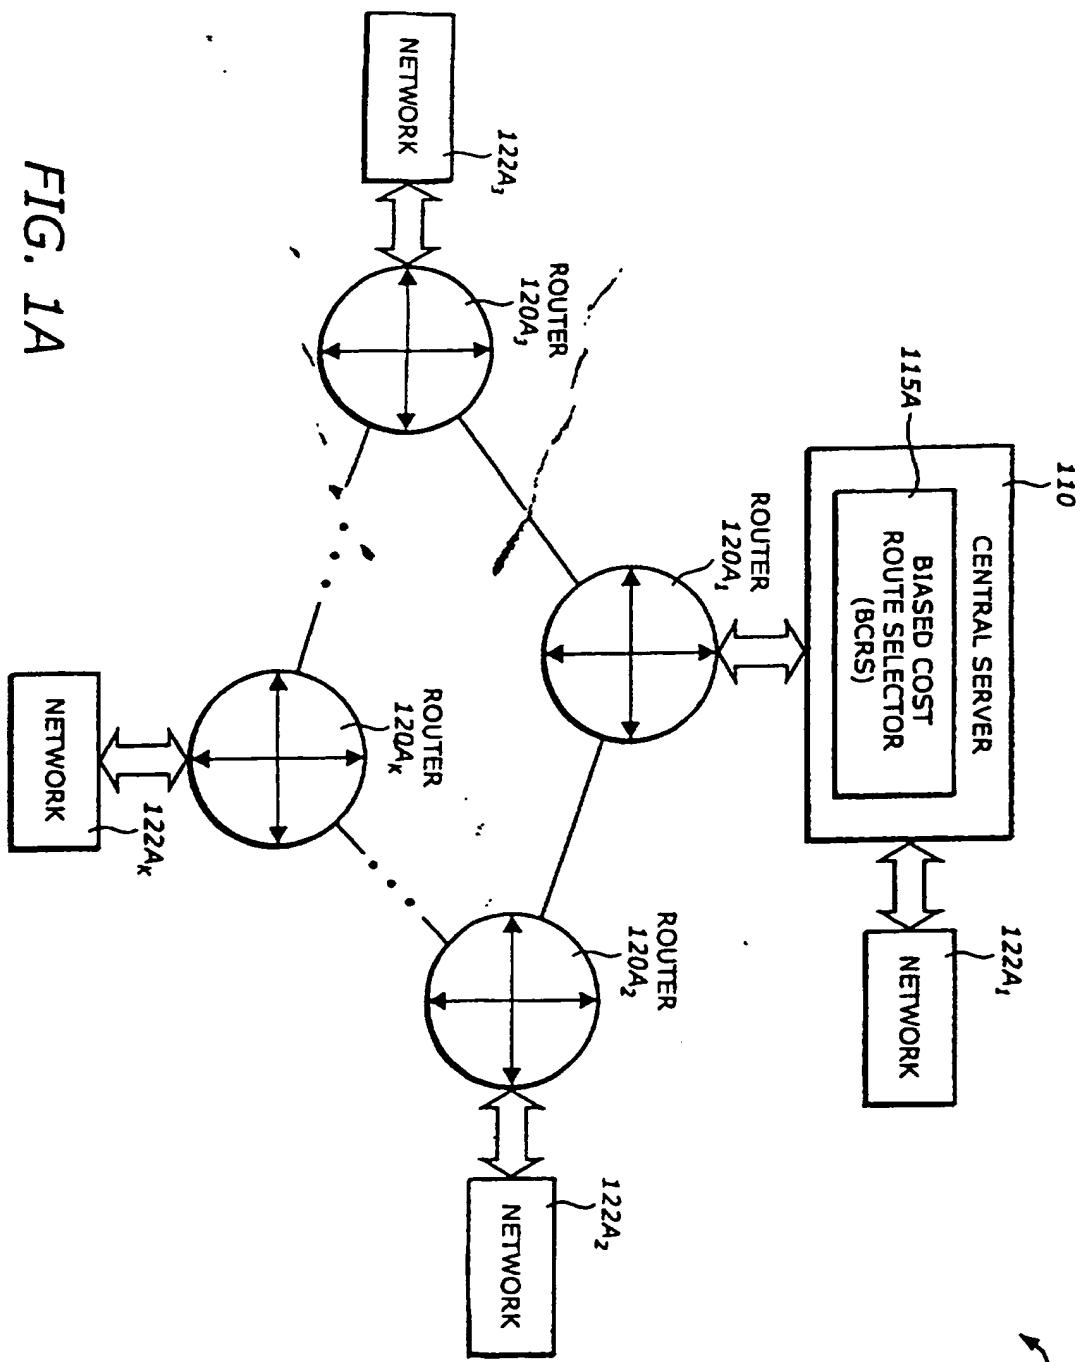
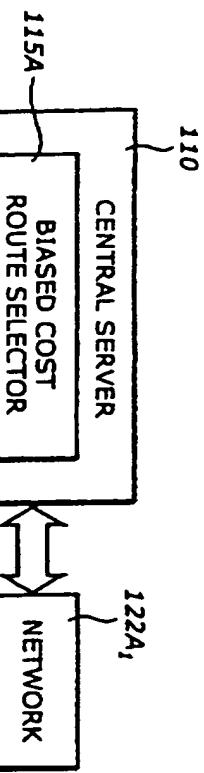


FIG. 1A

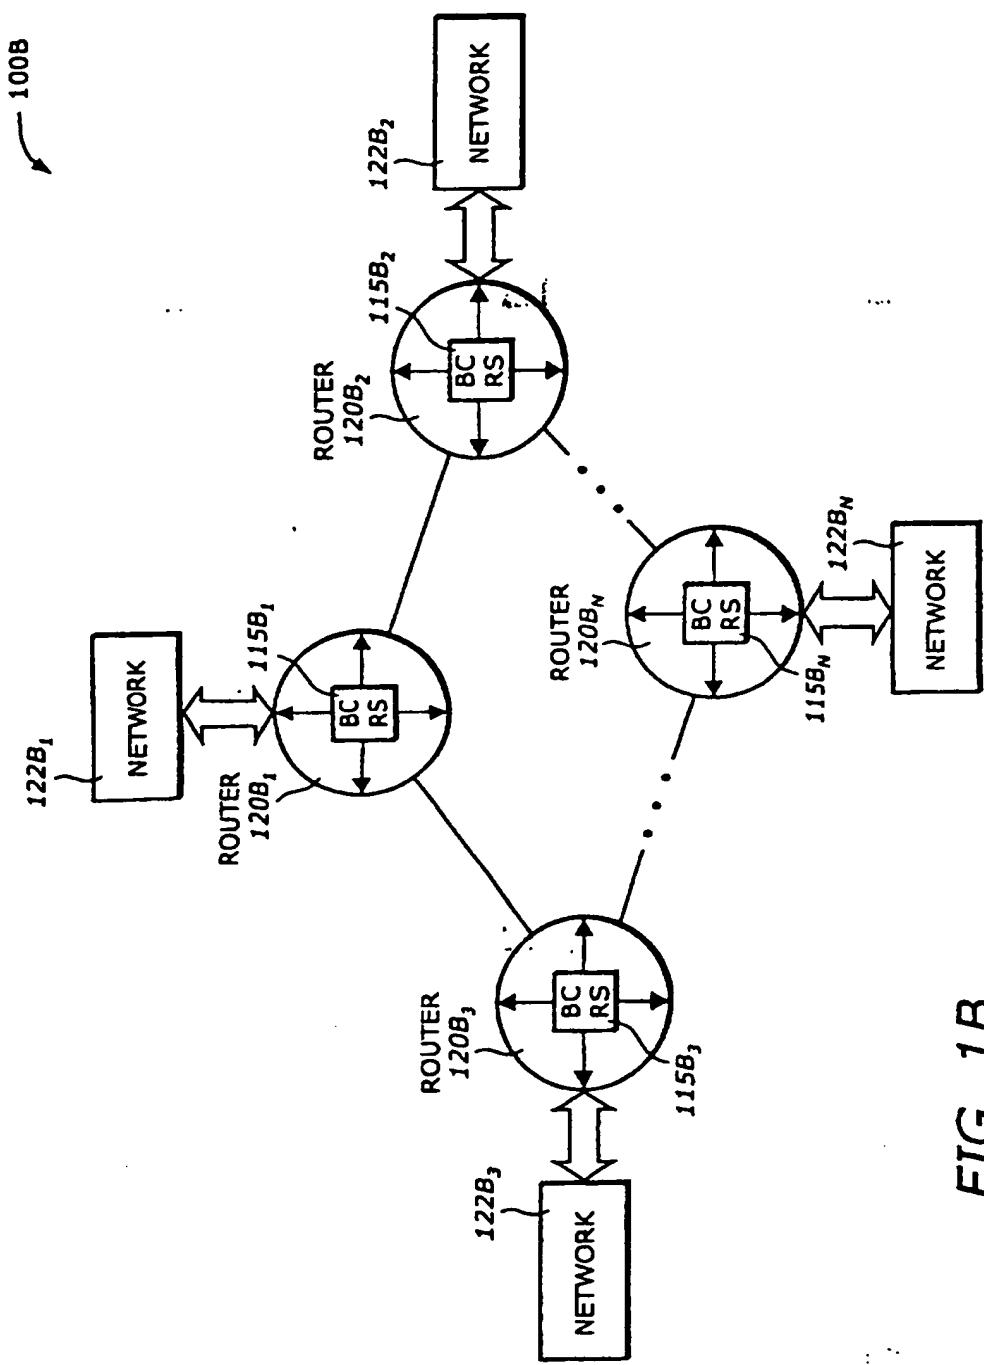


FIG. 1B

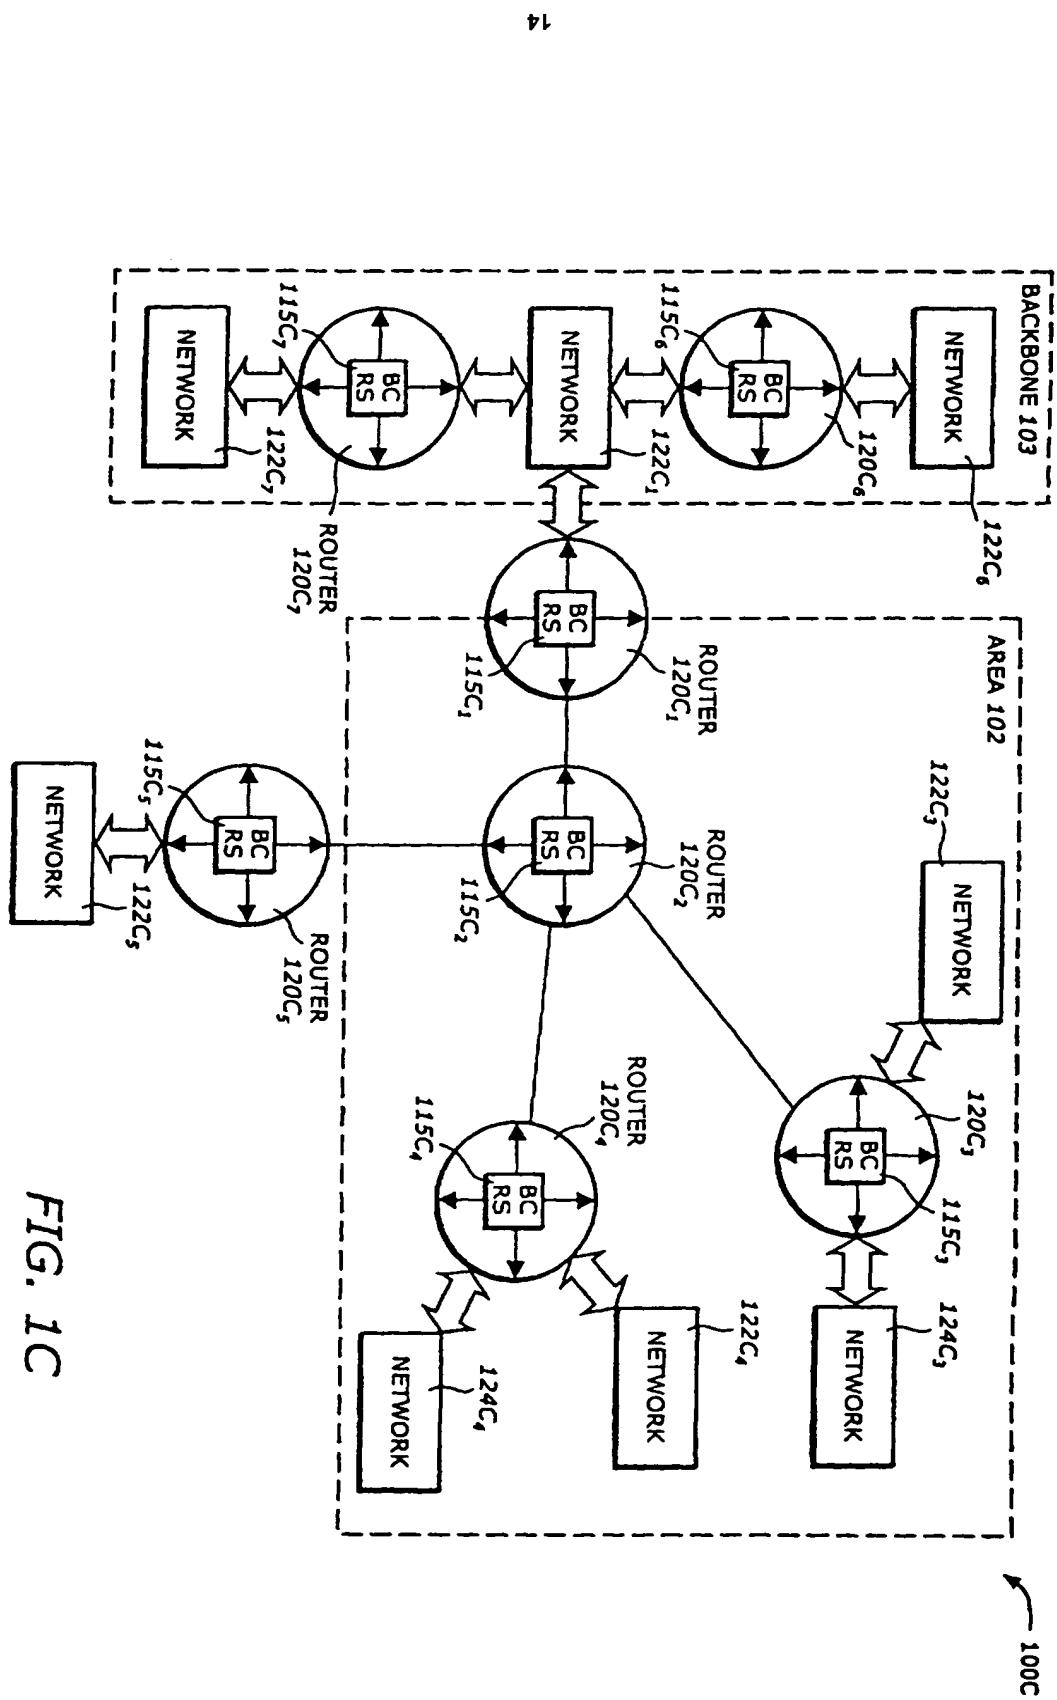


FIG. 1C

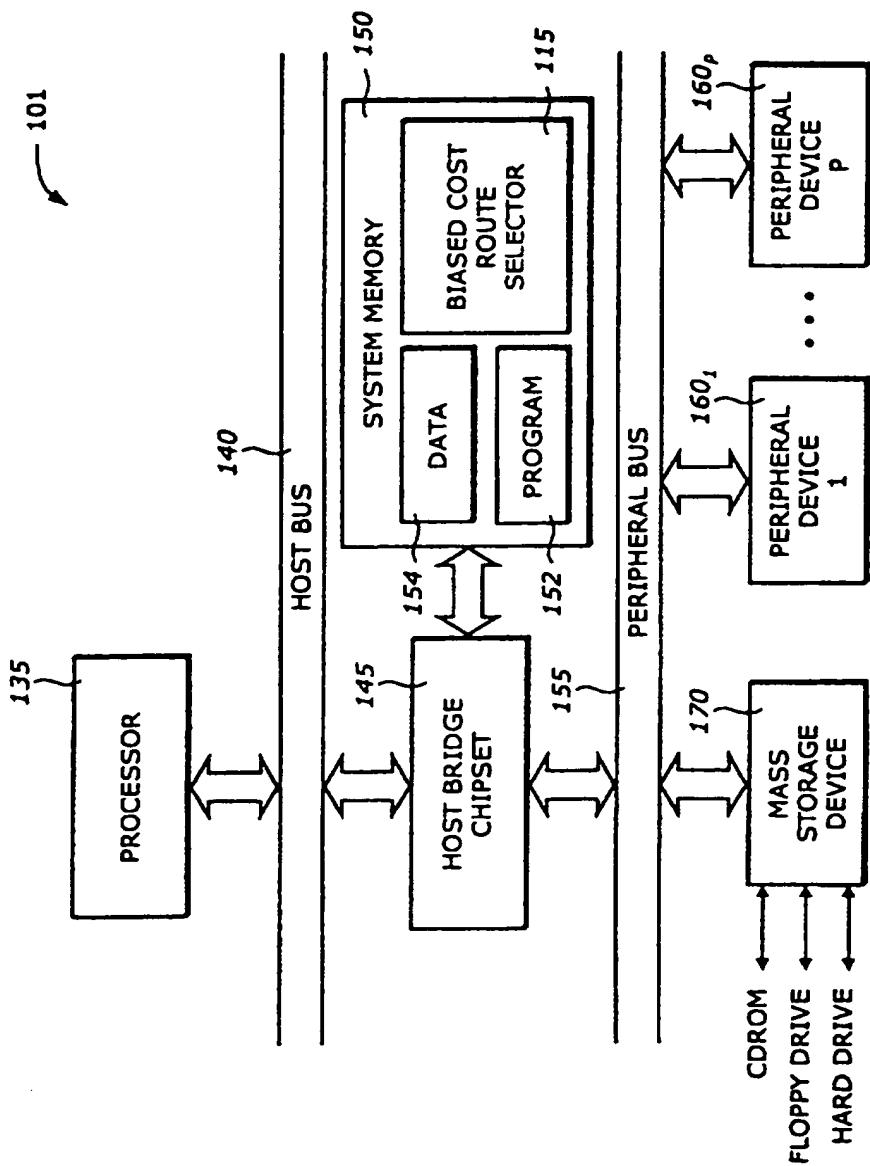
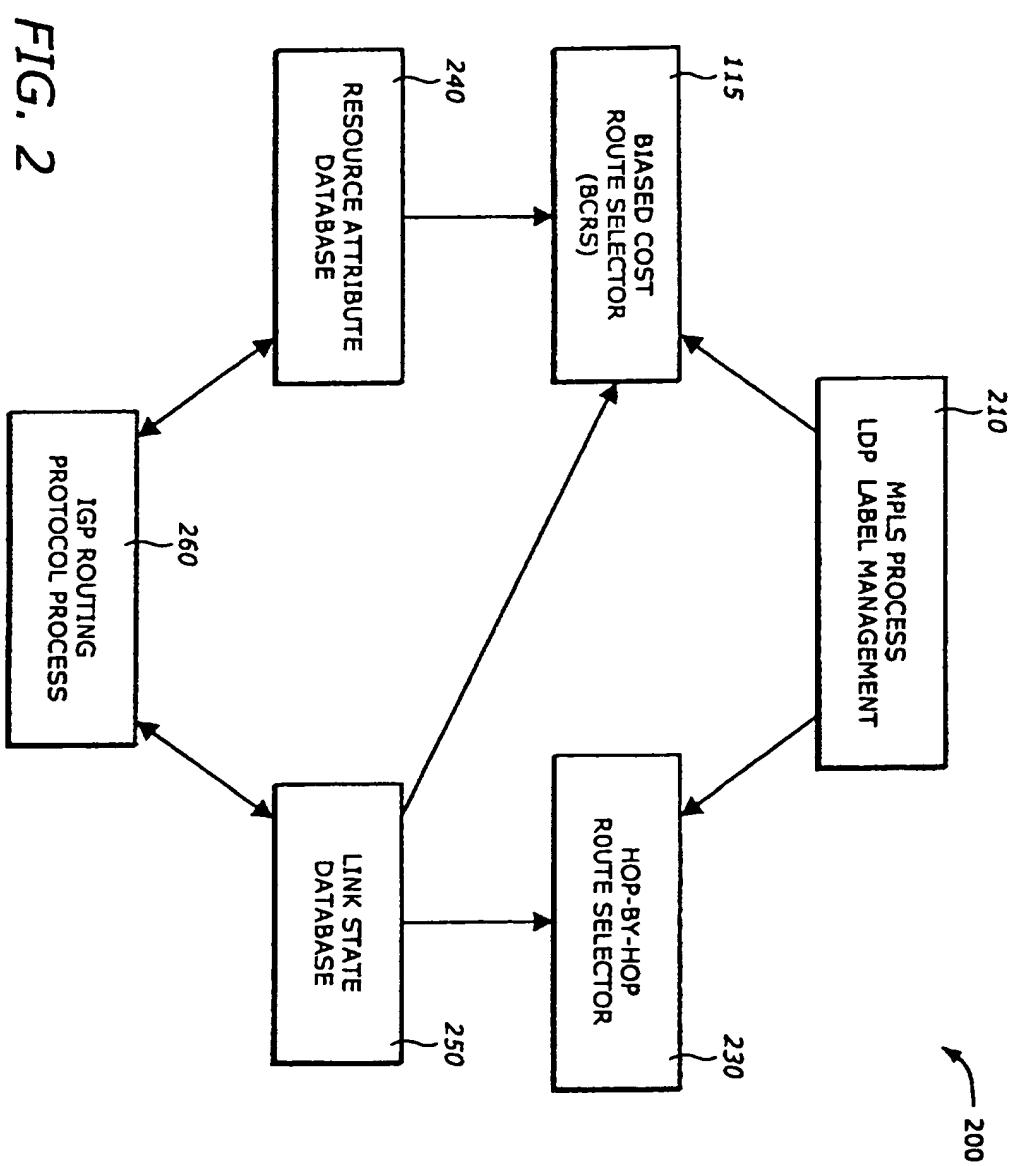


FIG. 1D



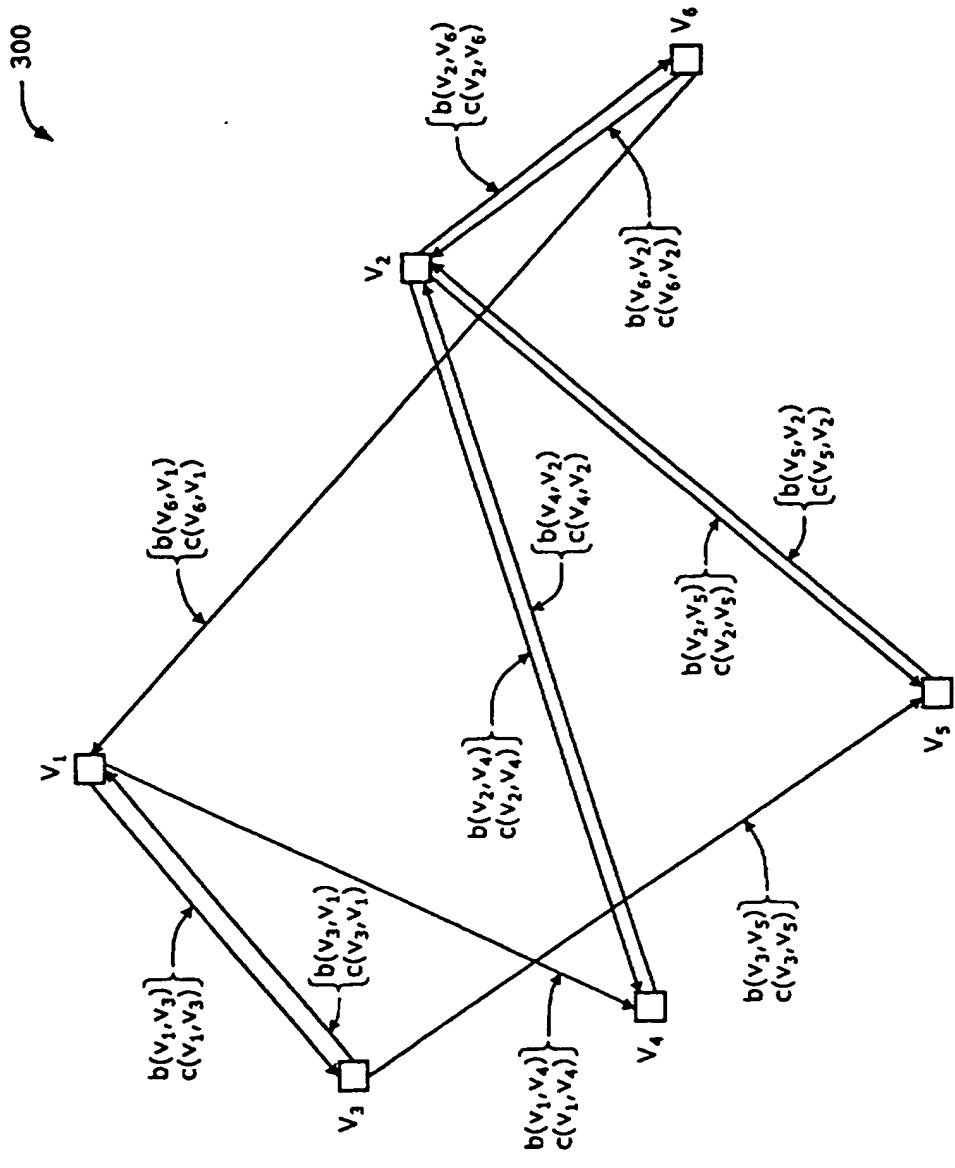
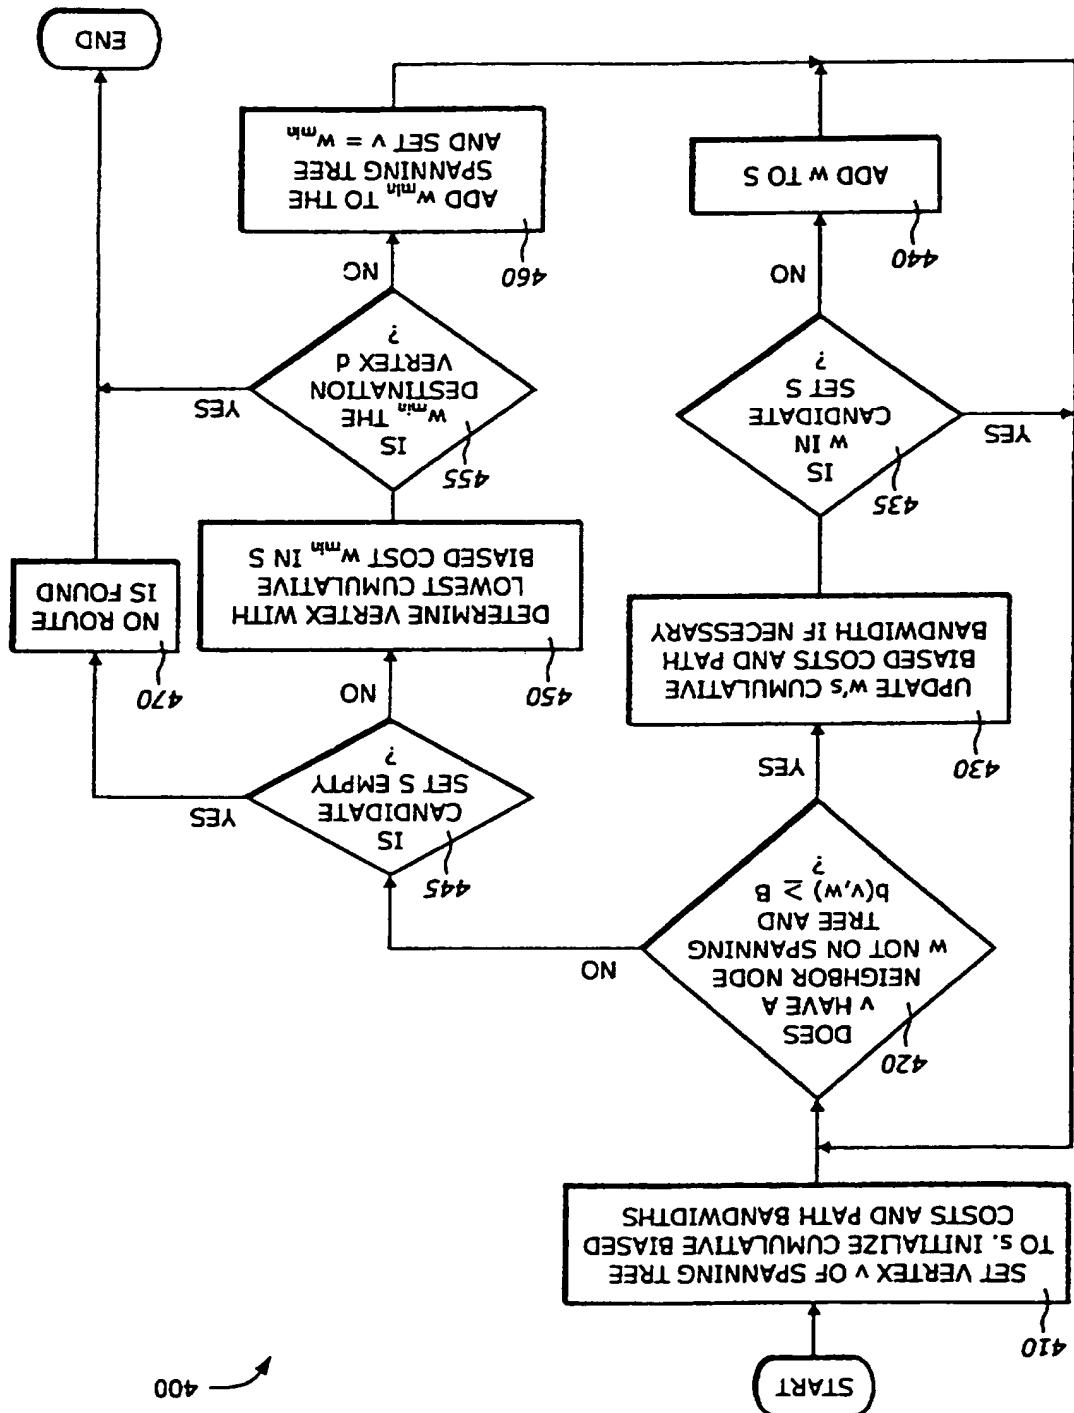


FIG. 3

FIG. 4



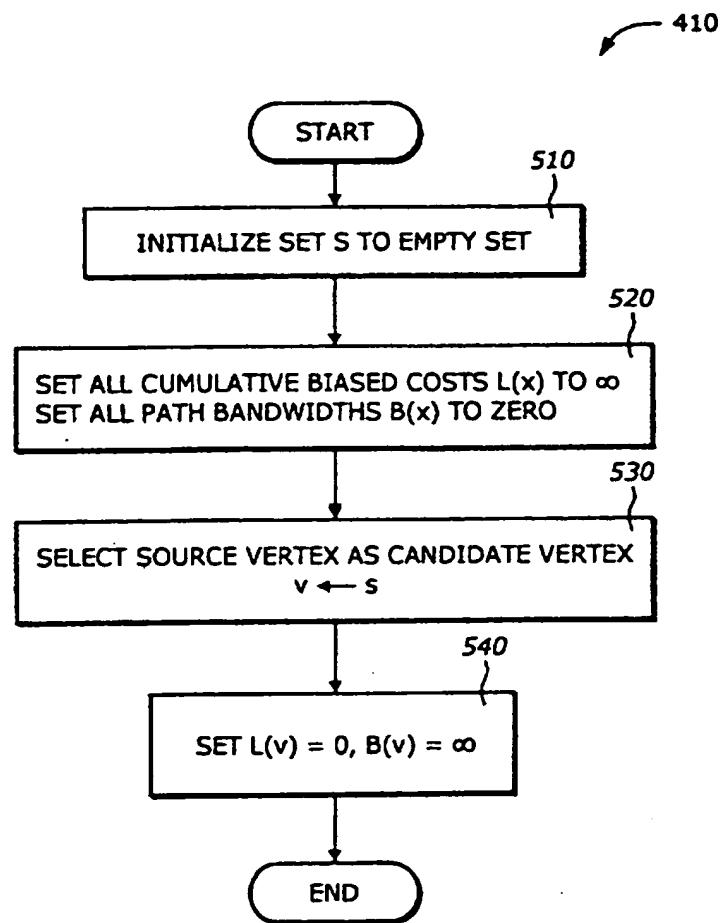
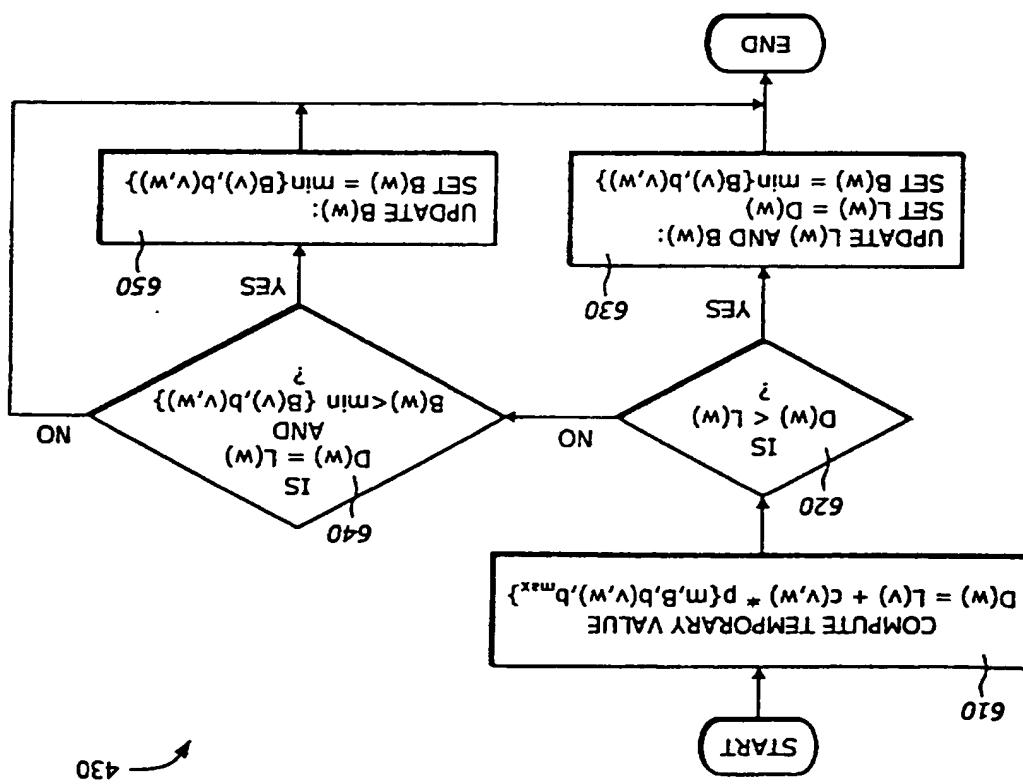


FIG. 5

FIG. 6



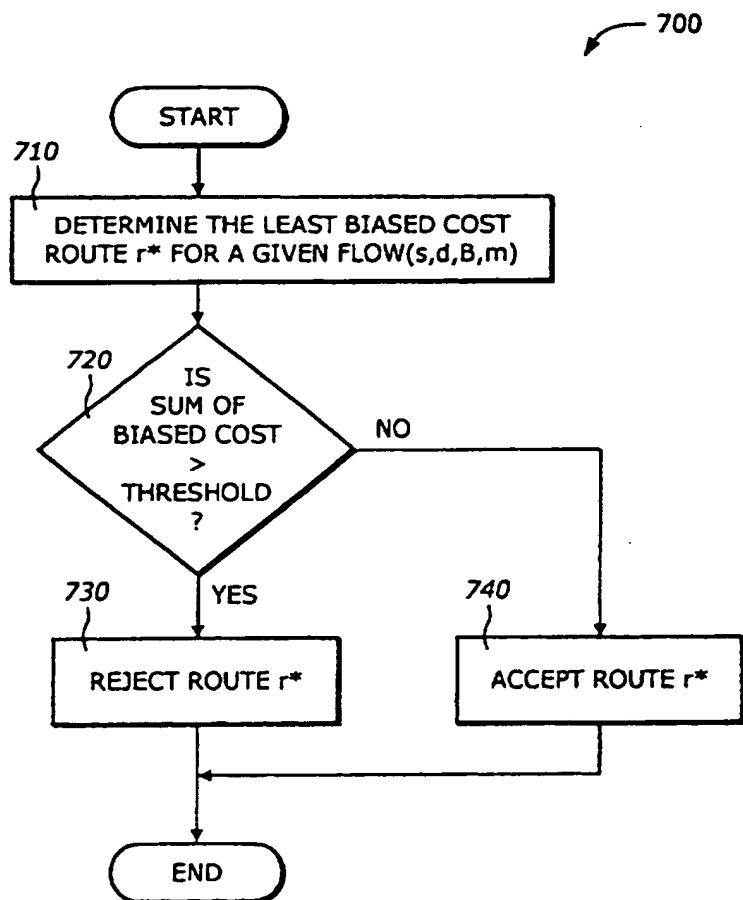
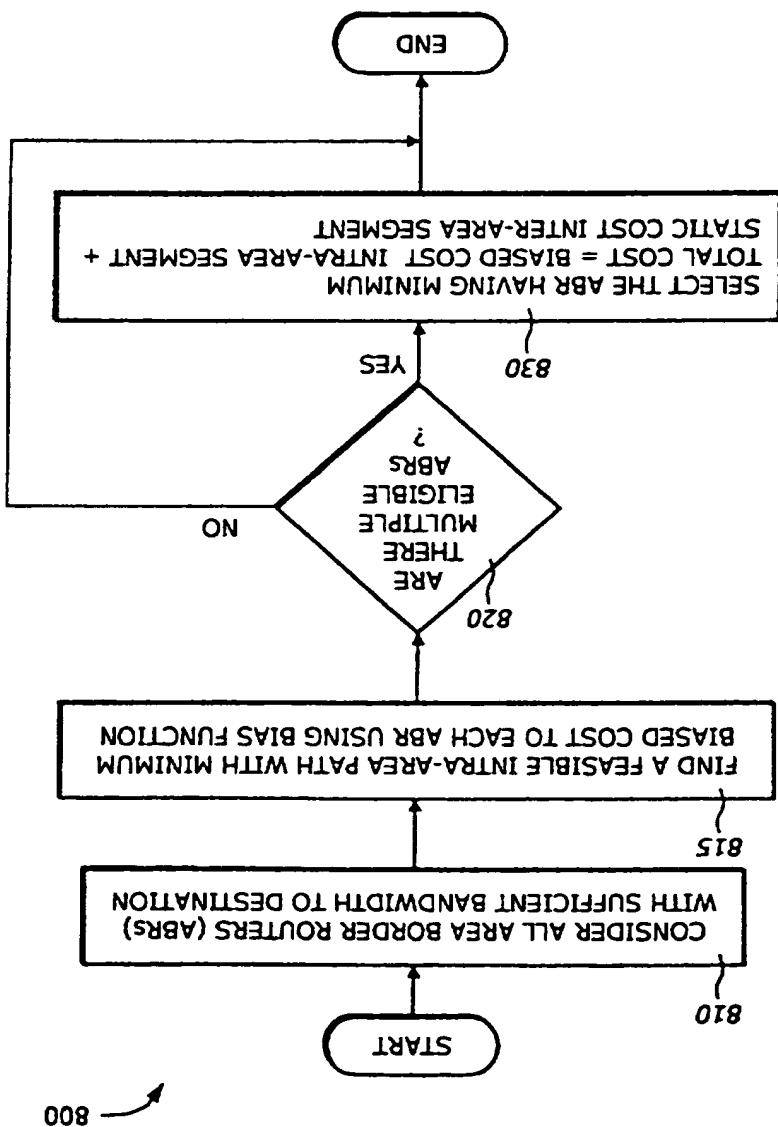


FIG. 7

FIG. 8A



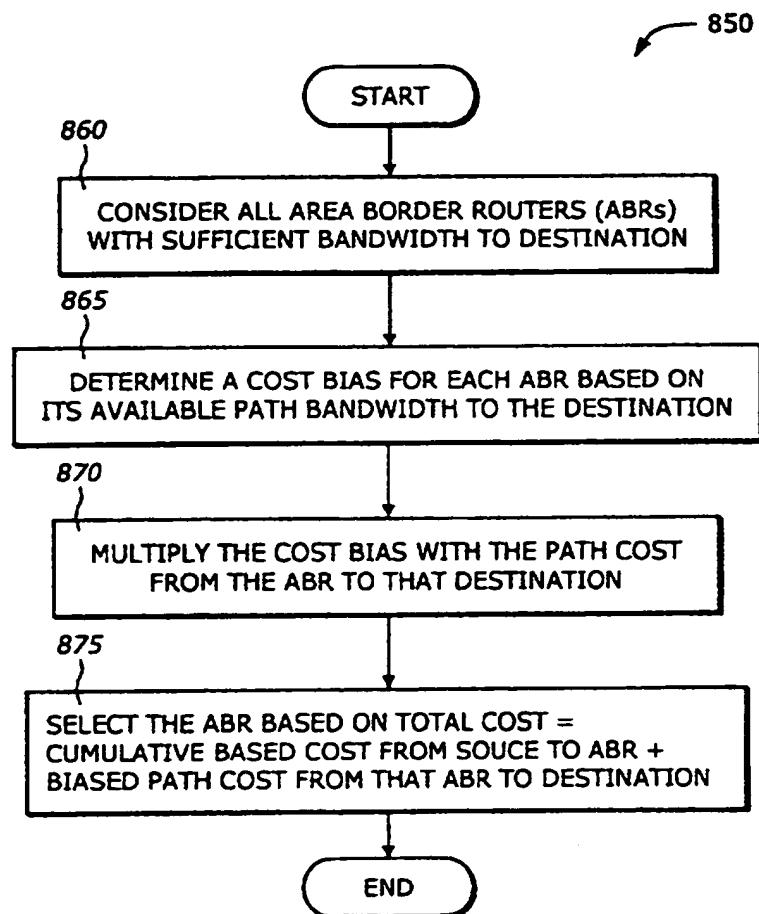


FIG. 8B

FIG. 9B

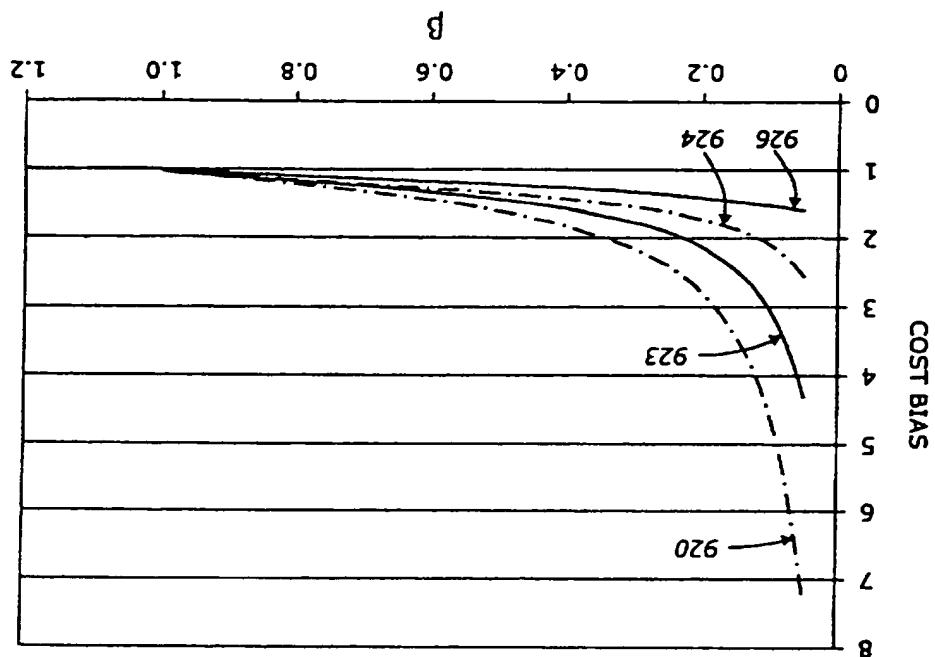
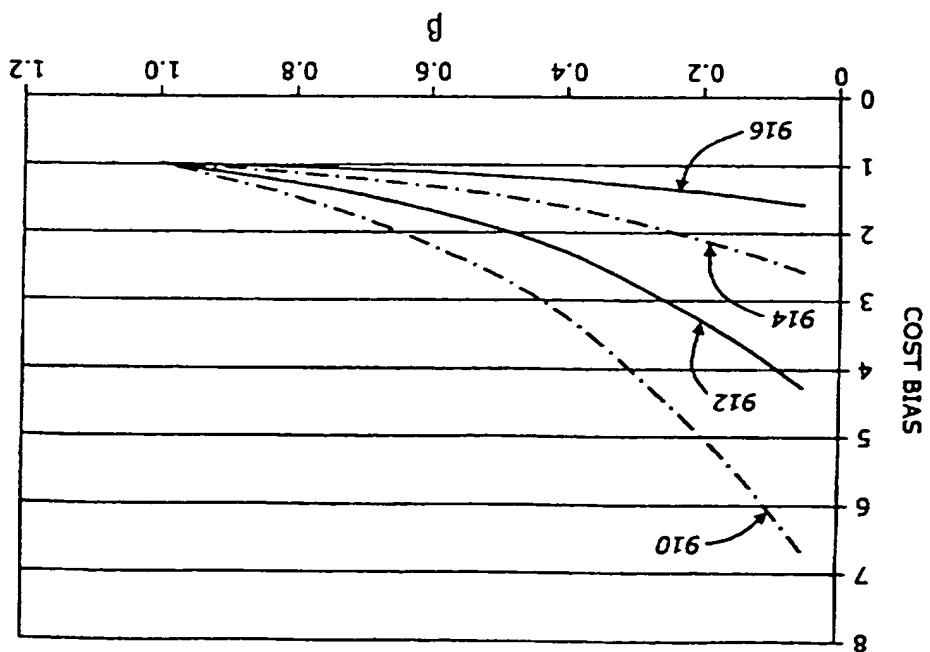


FIG. 9A



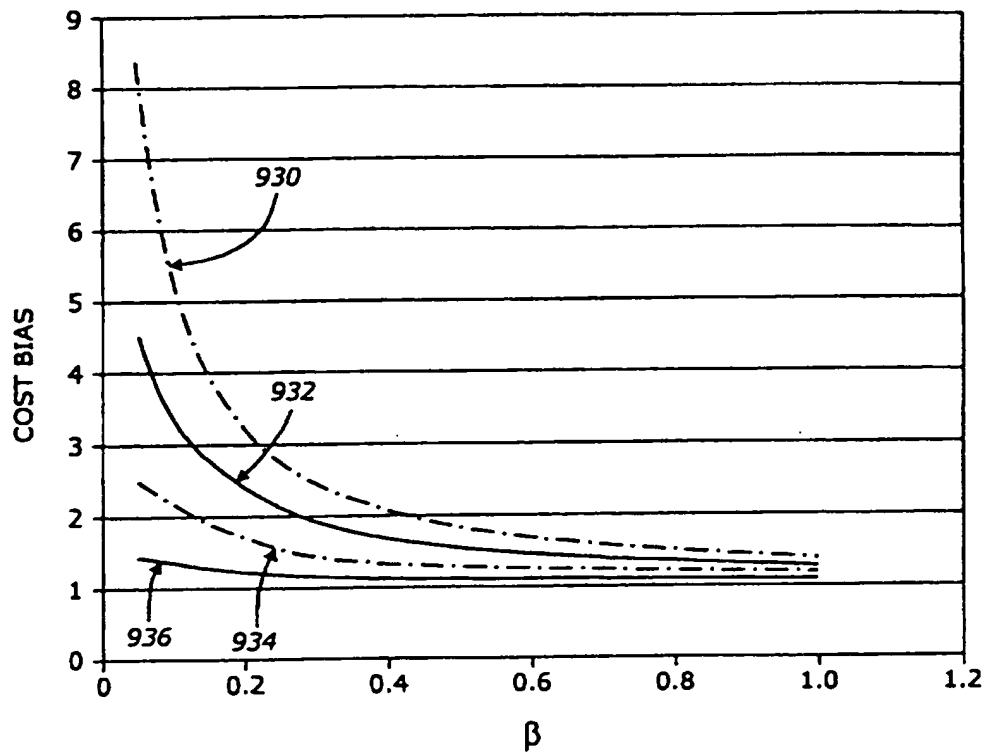


FIG. 9C

THIS PAGE BLANK (USPTO)